Do Privacy Controls Influence Content Generation and Sharing Patterns of Online Social Network Users? A Natural Experiment

Huseyin Cavusoglu
University of Texas at Dallas
huseyin@utdallas.edu

Tuan Phan
National University of Singapore
phantq@comp.nus.edu.sg

Hasan Cavusoglu
University of British Columbia
hasan.cavusoglu@sauder.ubc.ca

Introduction

In the age of social media, online social networks (OSNs), such as Facebook, Twitter, or LinkedIn, are indispensable for the majority of Internet users. To some, OSNs are as essential to their lives as electricity. These platforms enable OSN users to share content including status updates, pictures videos, comments, tags, and messages, with other people. Although most OSNs have been frequently criticized by privacy groups for the lack of attention to and care for privacy of their users, these platforms have continued to make inroads to new user bases. Today, OSNs are one of the fastest growing and most visited sites on the Internet. Facebook, the largest OSN, has reached 1 billion active users in September 2012 (Fowler 2012).

Despite the popularity of OSNs among Internet users, these platforms still suffer from a lack of a proven business model. Advertising is the main, and in most cases the only, source of revenue as users sign up for these services for free. Yet, OSNs seem far behind in capitalizing on their large user bases for advertising revenues compared to other web businesses. For instance, Facebook generated around $4 billion revenue last year, which is much less than the nearly $38 billion that Google earned from an equally large audience (Fowler 2012). While Google and Yahoo make about $88 from each user who uses its search engine, Facebook makes around $15 from each user (Fowler 2012).

The main advantage of OSNs over other web businesses, such as search engines, is their ability to access people and their networks of friends directly. Instead of tracking people's web activities and running complex data analytics tools to create user profiles in order to predict user interests, OSNs effortlessly gain insights into the lives of people through the content users voluntarily generate on their sites. This unique aspect creates a very powerful advertising tool not only in matching people with advertisers, but also in spreading the message across the network of friends, thereby facilitating social influence (Aral and Walker 2011). However, this opportunity extensively hinges on OSN

users creating more content and sharing it openly on the platform and engaging more with each other. Therefore, majority of OSNs have aimed at promoting content generation and openness in sharing content among their users, with an understanding that has been long been recognized by the social marketers: "*Encourage people to be public, increase ad revenue*" (Naone 2010). In explaining Facebook's interest in openness, Barry Schnitt, director of corporate communications and public policy, said that "*becoming less private and more public is a change just like it was a change in 2006 when Facebook became more than just people from colleges. Facebook is changing, and so is the world*." He also said "*By making the world more open and connected, we're expanding understanding between people…From a business perspective, if users are finding more value from the site, they will come back more and engage in more activity. And you can imagine the business consequences of that.*" (Kirkpatrick 2009) As alluded by Schnitt, increased openness, i.e., being more public and less private in sharing content, facilitates not only more traffic to the OSN platforms, but also more advertising revenues as OSNs can better understand the users when they divulge things about their lives and their favorite things, including products, music, movies and more. OSNs can measure what people are talking about and leverage it for real-time search, giving advertisers an accurate picture of users' interests and helping them deliver relevant ads to target audiences (Williams 2012).

In this paper, we study the relationship between privacy controls and content sharing patterns of users in the context of a specific OSN, hereafter referred to as the OSN platform. Although the OSN platform has revamped its privacy policy several times over the years, the changes were mostly dealing with default settings regarding different classes of personal information in user profiles (Freeman 2012). With these changes, the OSN platform aimed to make user profiles and generated content widely accessible via permissive default settings. Most of these changes have not even rolled out a new privacy control but simplified the notoriously complex privacy settings. However, these changes have been heavily criticized by privacy groups because they were seen as attempts to erode the privacy of OSN users and gradually lower the expectation of privacy (Anderson 2010). After studying the changes made by the platform in the past, the Electronic Frontier Foundation concluded that "*Viewed together, the successive policies tell a clear story. [OSN platform] originally earned its core base of users by offering them simple and powerful controls over their personal information. As [OSN platform] grew larger and became more important, it could have chosen to maintain or improve those controls. Instead, it's slowly but surely helped itself — and its advertising and business partners — to more and more of its users' information, while* ***limiting the users' options to control their own information***" (Opshal 2010). Control is obviously an issue when privacy is concerned. To address the concerns of privacy advocates about

limited controls, the backlash of its users, and the possible regulation by the government, the OSN platform has made a major change in privacy protection in December 2009. On December 9, the OSN announced that it has revamped the tools for privacy to enable its users to better control over information they share on the site (Sanghvi 2009). Apart from a simple privacy settings page, this change made it possible for users to apply privacy controls to determine access permissions at a higher level of granularity than choosing the same audience for each post (Cheng 2009). For the first time, OSN users were able to define the intended audience for each wall post separately, instead of being forced to use the same audience (Sanghvi 2009). Previously, OSN users had to go with a universal audience selector – everyone, only friends, or friends of friends – for each and every post. With the change, at each posting instance, OSN users were able to not only choose a different audience, but also customize the audience to a specific group such as individual friends or user-created lists. More than 350 million users were asked to review and update their privacy settings. With this major change, the OSN aimed to respond to privacy critics in a big way. However, the change was heavily criticized by privacy groups again because the true intention was seen as not to give people more control to protect their privacy, but to nudge people to publicly share even more (Bankston 2009; Kincaid 2009). It was regarded that the OSN platform was pushing for more openness in sharing indirectly, this time by empowering users with a feeling of control rather than relaxing default access rights directly, as done in the past (Kincaid 2009). Even the OSN platform spokesperson alluded to this by saying *"…so long as they [users] feel in control of who sees what, everyone seeing they post will likely to be a good for most people*" (Kirkpatrick 2009).

Since the policy change was purely exogenous in that users did not expect the event, it provides a natural experiment setting to study the impact of giving users more privacy control over information they generate and share on the platform. Specifically, we are seeking to answer if this change, which introduced granular privacy controls, structurally affected the content sharing patterns of users, as supposedly intended by the OSN platform and worried by privacy groups. Has this change really increased the openness of disclosure patterns?

Using a panel data obtained from the OSN, we empirically test unobserved effects models to assess the impacts of the policy change on content disclosure patterns of OSN users. Our initial findings suggest that the introduction of granular privacy controls indeed impacted users' information revelation structure on the social networking site. The OSN users appear to behave less closed and/or more open in their sharing content. Specifically, OSN users generate fewer messages and

more wall posts after the privacy change. The privacy change event is significant in explaining the disclosure patterns of users even after controlling for age and month fixed effects. Overall, our findings imply that perception of control over disclosure facilitated by the policy change makes users less private and/or more public in terms of content sharing patterns. We also find that although users with a large friend network are more closed in terms of disclosure patterns (i.e., privateness index increases with the network size) high socially connected users become less conservative and more open in their disclosures after the change. Hence, the policy change is instrumental in stopping the growth of closedness in shared content and facilitating the openness instead. Finally, we observe that both males and females are more socially active in term of content production, and more open in terms of disclosure patterns, compared to users who do not specify their gender. To the best of our knowledge, this is the first study on the impact of a major privacy change on the content generation and sharing patterns of OSN users, and one of the first large-scale studies on the impact of a major privacy policy change in a major OSN platform on user privacy-related outcomes using non-survey based methods or lab experiments.

Relevant Literature

Online social network facilitate creation of rich personal profiles. A profile is "a representation of their [selves] (and, often, of their own social networks) - to others to peruse, with the intention of contacting or being contacted by others" (Gross and Acquisti 2005). In addition to contact information, the elements of profile data range from relatively innocuous (such as favorite music or book) fields to potentially sensitive (such as sexual orientation or political affinity) ones. Early studies in online social networks examined privacy behavior in relationship to profile visibility. The level of visibility of various pieces of profile data, or of the overall profile, was seen equivalent to the level of disclosure.

Given that Facebook and Myspace were the most prominent online social network sites - especially among college students, early privacy studies on OSNs almost exclusively focused on the users of these sites. However, there were some major differences in default visibility settings. Facebook was partitioned into "networks", each representing a specific college. The default setting was that only people in your college (network) could see your full profile (i.e., a profile was visible to "everyone" in your network), while all other users could see only your profile picture, name of your network, and name you provided in your profile. Because of the demarcation of the site into networks, Facebook was often referred to as a "walled garden" (Tufekci 2008). Myspace, on the other hand,

was open to everyone by default, therefore regarded as less private. In addition, no networks or subgroups were available. However, both sites allowed their users to restrict default permissive privacy (visibility) settings to "friends only", meaning that only users who are designated as a "friend" can access one's profile. Since the default visibility was "everyone" in case of Myspace and "everyone in the network" in case of Facebook, researchers downloaded and examined user profiles to understand if disclosure behavior restricts the audience to "friends only" or if users choose not to reveal, thereby withhold, information in some fields on their profile. Apart from field data, researchers also surveyed users (mainly college students) to infer the amount of information revealed and the usage of privacy settings, and to compare stated privacy attitudes with actual disclosure behavior. These studies collectively documented empirical evidence of widespread disclosure practices in early days. Jones and Soltren (2005) found that more than half of the students disclosed information about their favorite books, music, and interests, but much less (17.1%) disclosed their phone numbers. Stutzman (2006) concluded that students overwhelmingly disclosed their birthday, relationship status, and political view, while disclosure of cell phone number was limited to 16.4%. In addition to confirming the finding in other studies about high levels of disclosure of personal information on user profiles, Gross and Acquisti (2005) found that only a small set of users adjust the default (permissive) privacy settings to restrict the visibility of their profiles. Acquisti and Gross (2006) identified that stated privacy concerns have little influence on information disclosure behavior: highly concerned users also reveal extensive information on their profiles. Taken together these findings pointed out the dichotomy between stated privacy concerns and actual information sharing behavior. Lampe et al. (2007) reported that only 19% of profiles are set as "friends only". A survey study by Tufekci (2008) revealed that Facebook and Myspace users do not set their profile visibility in relationship to the level of their general privacy concern, but the fear of unwanted audience has an impact on profile visibility settings. However, perceived future audience (romantic partner, employee, government) has no impact on the visibility of their profiles, concluding that although users are better at managing "spatial" boundaries by restricting the visibility of their profiles to current audiences, they are less concerned about, or aware of, intrusions through "temporal" boundaries by future audiences. Similar disclosure practices were found among Myspace users as well (Caverlee and Webb 2008; Thelwall 2008). Different from other studies exclusively focusing on one's own profile attributes on selected privacy settings, Lewis et al. (2008) examined relational data – friendship and roommate ties- as factors that contribute to a student having a private profile instead of a public one. They defined a profile to be private if the student has

changed the default settings in a way that the profile is not fully accessible or searchable by non-friends in the same network, indicating privacy preserving behavior limiting visibility. They found that a student is more likely to have a private profile if (i) the student is female; (ii) the student's friends–especially roommates- have private profiles; (iii) the student is more active on Facebook; and (iv) the student prefer music that is relatively popular, and only relatively popular music.

With increased awareness of privacy threats and extensive coverage of these issues in popular press, OSN users started exhibiting more privacy seeking behavior over time. Stutzman et al. (2012) documented this evolution of privacy and disclosure behavior. Examining the profiles of early Facebook adopters in the CMU network, they concluded that users have reduced the amount of personal information on their profiles shared with other unconnected users in the same network between years 2005 and 2009. They also showed the reversal of the privacy seeking trend after 2009 as users have resumed to public sharing of various elements of profile data. Similar to our study, Stutzman et al. (2012) takes a longitudinal perspective. However, our study is different from theirs in several key aspects. First, they consider how disclosure of personal information in user profiles has changed since the early days of Facebook. Therefore, their focus is on profile elements, such as birthdate, political affiliation, home address, and others, not on content generating user activities, like wall posts or private messages that we study. We are exclusively focusing on content generation activities and analyzing the changing patterns in response to a specific policy change. Second, their study is mostly descriptive and argues whether public disclosure of personal profile information has been influenced by the changes in the default privacy settings of the OSN platform. On the other hand, we build econometric models to study the causal link between privacy controls enabled by the policy change and user generated contents and their patterns. Third, while they consider whether users reveal or withhold information in their profiles based on observations in different points in time (specifically, in seven yearly snapshots), we utilize weekly-level content generating activities across 192 weeks.

Prior research has examined how privacy controls influence disclosure patterns of OSN users and provided some evidence on the paradoxical effects. Using survey-based experiments, Brandimarte et al (2010) studied the relationship between perceived control over the release of personal information and individuals' propensity to disclosure this information. They showed that people indeed reveal more (less) when they have more (less) control over the information that they release. Similar to Brandimarte et al (2010), among other arguments, we attribute extensive information

disclosure behavior of OSN users after the policy change to the sense of control hypothesis. However, unlike Brandimarte et al (2010), we use field data to assess the impact of a real privacy change on dynamic content generation and sharing behavior, not just on (mostly) static profile data. In addition, we show the impact of increased control over release on disclosure using a panel dataset that spans around four years. Furthermore, we capture the sharing patterns of different types of user-generated contents instead of whether a user merely chooses to reveal profile data. Hence, we assess disclosure patterns, and therefore privacy concerns, in terms of the openness and closedness of content sharing rather than willingness to disclose personal information on profiles.

Theoretical Framework

Privacy is often defined as "the selective control of access to the self" in the literature (Altman 1975). Control and ownership of private information are interrelated concepts and salient aspects of privacy management. Controls determine who has access to information that belongs to us, preventing unwanted exposure. The Communication Privacy Management (CPM) theory suggests that people conceal or reveal their private information by coordinating interpersonal boundaries (Petronio 2002). While thick boundaries maintain high level of controls to promote secrecy, thin boundaries deploy fewer controls to facilitate openness. If information is private, one has a strong control over her information, and subsequently the boundary becomes very tight. One can choose to share her private information with one or more people. Those who are privy to her information become part of her cognitive information space with a clearly defined boundary. When the boundary is lax (tight), more (less) people are within the information space. This requires individuals to perform a cognitive calculus in order to determine whether to reveal or conceal the information, and whom to reveal (Laufen and Wolfe 1977). In doing so, they consider the benefits, such as social capital, and the costs, such as the risk of misusing of private information (Petronio 2002, Margulis 2003). CPM theory highlights the dialectical tensions between openness and closeness. There is a need of being private through concealing, and also a need of being public through revealing (Petronio 2002).

Once a person discloses her private information, he/she becomes a little bit less private/more public. Therefore, an action toward disclosing is also an action against privacy. On the OSN platform, people always reveal information. We cannot observe what information people conceal. However, while sharing information with others, people can choose whom to share with. People can tightly control their boundaries by sending private messages, which reflects the *closeness in sharing.*

Alternatively, people can loosely control their boundaries by making (public) wall posts, which captures *openness in sharing*. Hence, *privateness index,* a proxy for disclosure pattern of a user, defined as the ratio of number of private messages to the number of wall posts in a given period represents how tight/loose the boundary around information that is shared. The higher the privateness index, the more private (or closed) the individual is in a given period in terms of content sharing patterns.

Theory of Planned Behavior (TPB) argues that an individual's perceived behavioural control, along with her attitude and subjective norm, influences her behavioural intention and actual behaviour (Ajzen 1991). Behaviour is partially driven by *perceived* controllability, which might not necessarily match to the *actual* controllability. As long as the individual perceives that she is in control of his/her behaviour, he/she is more likely to perform the behaviour. Furthermore, Peltzman (1975) argues that individuals exhibit a tendency to react to a safety regulation by increasing other risky behaviour, offsetting some or all of the benefit of the regulation. People also have a tendency to overestimate their ability to affect the outcome of events when they can exercise control (Thompson 1999). For instance, people generally feel more comfortable when they are in a driver seat (i.e., high-control situation) than in a passenger seat (i.e., low-control situation). This perception is more pronounced when people have choices in terms of controls (Langer 1975). One possible explanation for this phenomenon is that people desire to avoid the negative consequences associated with having no control over outcome. Fischhoff et al. (1978) found out that perceived level of "control" influences the relationship between risk perception and risk acceptance: the lack of control results in judging risks more severely whereas having controls leads to evaluate risks less severely than they actually are. On the OSN, if an individual perceives that he/she is better equipped with controls to allow him/her to manage his/her privacy; he/she is likely to loosen the boundary surrounding his/her private information, regardless of whether the user makes use of those features. We argue that the late 2009 privacy policy change in the OSN platform afforded users with more granular controls over disclosure of their private information, and subsequently affected the users' perceptions that they are more in control than before. Hence, consistent with theory of planned behaviour, perception of control and risk studies, we postulate that the this policy change will result in open information sharing, in which individuals become less private and/or more public in terms of their content disclosure patterns.

Dataset

We use backend data from a major OSN to assemble a panel data. We sampled data on 13,145 OSN users and their activities for 192 weeks, from September 3, 2007 (week 1) till May 23, 2011 (week 192), inclusive. The policy change event took place during week 119. We aggregated user activities at a weekly level. Our overall dataset contains 2,428,885 observations (user-week pairs). We considered two most popular content sharing activities: (i) wall posts and (ii) messages. Wall posts, including status updates, links, pictures and videos, by their nature, have an intended audience of more than one user (facilitating *openness* in sharing) whereas messages are typically directed at only one or few users (facilitating *closeness* in sharing). In addition, we have other covariates that may affect content generation and disclosure patterns of users. These variables are OSN Age (number of weeks since a user joined the OSN), gender (not specified, female, male), and network size (number of friends of a user before that week).

Analysis and Results

We analyze the potential impact of the policy change event on three dependent variables: (i) wall posts, (ii) messages, and (iii) privateness index. We call the total number of private messages the user *i* generates in week *t* as $Message_{it}$. Similarly, we call the total number of wall posts the user *i* generates in week *t* as $Post_{it}$. Using these two values, we define the ratio of $Message_{it}$ to $Post_{it}$ as privateness index. This index, which captures the information disclosure patterns of users, has a nice interpretation. The higher the value of this index, the higher (lower) the closedness (openness) of content sharing on the OSN platform. Although this relative index is very simple, it captures valuable information about the sharing patterns of different user-generated content activities. If a user generates more wall posts, the openness of content sharing will increase. Similarly, if a user generates more private messages, the openness of content sharing will decrease. If a user increases the numbers of wall posts and messages at the same rate, openness of his/her sharing pattern will not change (though the amount of generated content will increase).

Baseline Panel Models

In this section, we assess the influence of the privacy policy change OSN users and their sharing behavior. We seek answers to the following questions: Does the policy change have a lasting effect

on user disclosure patterns? If so, what kind of influence does the privacy policy change facilitate in terms of content sharing and privateness of this sharing?

To assess the impact of the OSN policy change, we build unobserved effects panel models as

$$Ln(Message_{it}) = \alpha_i^m + \beta_1^m AfterPolicy_t + \beta_2^m OSNAge_{it} + Month\ Fixed\ Effects + \varepsilon_{it}^m \quad (1)$$

$$Ln(Post_{it}) = \alpha_i^p + \beta_1^p AfterPolicy_t + \beta_2^p OSNAge_{it} + Month\ Fixed\ Effects + \varepsilon_{it}^p \quad (2)$$

$$Ln(\frac{Message_{it}+1}{Post_{it}+1}) = \alpha_i^r + \beta_1^r AfterPolicy_t + \beta_2^r OSNAge_{it} + Month\ Fixed\ Effects + \varepsilon_{it}^r \quad (3)$$

where α_i's capture the activity-specific unobserved heterogeneity of user *i*. $AfterPolicy_t$ is our privacy change event dummy, which takes the value of 1 if week *t* is after Dec 9, 2009, and zero otherwise. We use the natural logarithms of $Message_{it}$, $Post_{it}$, and their ratio as these variables are highly skewed[1]. We also include $OSNAge_{it}$ and $Month\ Fixed\ Effects$ as control variables. Finally, ε_{it}'s are idiosyncratic errors with standard assumptions. We estimate these models with both the fixed-effects (FE) estimator and the random-effects (RE) estimator.

Table 1. The Impact of the Privacy Policy Change in Information Disclosure Patterns

Dep. Variable	*Ln(Message)*		*Ln(Post)*		*Ln((Message+1)/(Post+1))*	
Model	(1a)	(1b)	(2a)	(2b)	(3a)	(3b)
AfterPolicy	-0.0134***	-0.0324***	0.0419***	0.0271***	-0.0553***	-0.0385***
	(0.0015)	(0.0015)	(0.0013)	(0.0013)	(0.0017)	(0.0017)
OSNAge	-0.0011***	-0.0009***	-0.0023***	-0.0021***	0.0012***	0.0010***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Month Effects	Included	Included	Included	Included	included	included
Constant	0.4570***	0.4290***	0.6830***	0.6540***	-0.2260***	-0.2070***
	(0.0018)	(0.0037)	(0.0015)	(0.0042)	(0.0020)	(0.0034)
Specification	FE	RE	FE	RE	FE	RE

The sample includes 2,428,885 observations from 13,145 users.

Robust standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

[1] We add one to the numbers of weekly messages and posts before taking the logarithms as some numbers can be zero.

We can observe from the estimation results in Table 1 that there is a clear pattern in the results irrespective of the estimation method used. The negative and significant coefficient of *AfterPolicy* in models (1a) and (1b) indicates that users generate fewer messages after the privacy change. The positive and significant coefficient of *AfterPolicy* in models (2a) and (2b) implies that users generate more posts after the privacy change. In addition, consistent with the results regarding the opposing effects of the policy change on the number of messages and posts, we can see that OSN users become more open in terms of content sharing patterns. The coefficient of *AfterPolicy* in models (3a) and (3b) is negative and significant. Furthermore, as user age increases, we can see that both types of content sharing activities drop. However, the reduction is more pronounced in wall posts, and therefore privateness index increases with tenure in the platform. All these results are in line with our expectations. It seems that the change has a significant influence on content sharing activities and the resulting disclosure patterns of OSN users, as intended by the platform and as predicted by the theory. That is, in reactions to the change, users reduce the number of messages and increase the number of wall posts, which in turn, result in a drop in the closeness in sharing (or increase in the openness in sharing).

Impact of Network Size and Gender in Disclosure Behavior

It is possible that users with varying network sizes can have different response to the privacy change because network size has implications for the reach of wall posts of a user. In general, users with more friends have a larger audience for their wall posts. Apart from network size, gender can be an important factor in capturing the reaction to the privacy policy change. It has been shown in prior studies that females are different from males in their sensitivity and therefore activities in online social networks. We examine the impact of network size and gender on disclosure behavior. We operationalize the network size with the number of friends of a user *i* in week *t*, $Friend_{it}$. We use gender dummies, $Male_i$ and $Female_i$.[2] We also add the interactions of network size and genders with the privacy event. Specifically, we extend our baseline models and estimate the following extended models.

[2] Unlike today, specifying a gender to sign up for the OSN service was not required in the past. Therefore, we have some users who do not have any gender information. We used them as a baseline.

$$Ln(Message_{it}) = \alpha_i^m + \beta_1^m AfterPolicy_t + \beta_2^m OSNAge_{it} + \beta_3^m Friend_{it} + \beta_4^m Male_i + \beta_5^m Female_i + \beta_6^m Friend_{it} * AfterPolicy_t + \beta_7^m Male_i * AfterPolicy_t + \beta_8^m Female_i * AfterPolicy_t + Month\ Fixed\ Effects + \varepsilon_{it}^m \quad (4)$$

$$Ln(Post_{it}) = \alpha_i^p + \beta_1^p AfterPolicy_t + \beta_2^p OSNAge_{it} + \beta_3^p Friend_{it} + \beta_4^p Male_i + \beta_5^p Female_i + \beta_6^p Friend_{it} * AfterPolicy_t + \beta_7^p Male_i * AfterPolicy_t + \beta_8^p Female_i * AfterPolicy_t + Month\ Fixed\ Effects + \varepsilon_{it}^p \quad (5)$$

$$Ln(\frac{Message_{it}+1}{Post_{it}+1}) = \alpha_i^r + \beta_1^r AfterPolicy_t + \beta_2^r OSNAge_{it} + \beta_3^r Friend_{it} + \beta_4^r Male_i + \beta_5^r Female_i + \beta_6^r Friend_{it} * AfterPolicy_t + \beta_7^r Male_i * AfterPolicy_t + \beta_8^r Female_i * AfterPolicy_t + Month\ Fixed\ Effects + \varepsilon_{it}^r \quad (6)$$

Table 2 presents our estimation results. The negative and significant coefficient of *Friend*AfterPolicy* in models (4a,4b) and (5a,5b) indicate that highly connected users are affected more negatively and contribute less to online social networking activities after the change. That is, OSN users with more friends reduce the number of messages and wall posts more than users with fewer friends. In addition, although users with a large friend network are more closed in terms of disclosure patterns (i.e., privateness index increases with the network size) the negative coefficients of the interaction term in models (6a, 6b) suggest that socially connected users become less conservative and more open in their disclosures after the change. Assuming that network size of an average user increases over time, an increase in privateness index over network size implies that the OSN platform fails to motivate its users for continuously open engagement as they get older. However, the privacy change effectively reverses this trend. That is, the policy change is instrumental in stopping the growth of closedness in shared content and facilitating the openness instead. As for the influence of gender, we can observe that both males and females are more socially active in term of content production, and more open in terms of disclosure patterns, compared to users who do not specify their gender. This is expected given that users who choose not to reveal their gender may, in general, be more privacy conscious. Interestingly, although the privacy change increases the openness of everyone, users who reveal their gender as males and females reduce their openness and become more private in their disclosures.

Table 2. The Impact of the Privacy Change on Users with Different Network Size and Gender

Dep. Variable	*Ln(Message)*		*Ln(Post)*		*Ln((Message+1)/(Post+1))*	
Model	(4a)	(4b)	(5a)	(5b)	(6a)	(6b)
AfterPolicy	0.0707***	0.0906***	0.1420***	0.1600***	-0.0716***	-0.0794***
	(0.0018)	(0.0018)	(0.0015)	(0.0015)	(0.0021)	(0.0020)
Friend	-0.0002***	0.0002***	-0.0008***	-0.0004***	0.0006***	0.0003***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
*Friend*AfterPolicy*	-0.0007***	-0.0008***	-0.0004***	-0.0005***	-0.0003***	-0.0002***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Male		0.5780***		0.8500***		-0.2410***
		(0.0066)		(0.0070)		(0.0067)
Female		0.6680***		1.0320***		-0.3330***
		(0.0065)		(0.0069)		(0.0065)
*Male*AfterPolicy*	0.0454***	0.0332***	-0.0986***	-0.1100***	0.1440***	0.1510***
	(0.0021)	(0.0022)	(0.0018)	(0.0018)	(0.0025)	(0.0025)
*Female*AfterPolicy*	-0.0360***	-0.0494***	-0.1140***	-0.1270***	0.0783***	0.0858***
	(0.0021)	(0.0021)	(0.0018)	(0.0018)	(0.0025)	(0.0025)
OSNAge	-0.0008***	-0.0010***	-0.0016***	-0.0018***	0.0008***	0.0009***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Month Effects	included	included	included	included	included	included
Constant	0.4490***	0.0748***	0.6930***	0.1370***	-0.2440***	-0.0666***
	(0.0018)	(0.0042)	(0.0015)	(0.0044)	(0.0021)	(0.0043)
Specification	FE	RE	FE	RE	FE	RE

The sample includes 2,428,885 observations from 13,145 users.
Robust standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Ongoing Research

Although our initial results show that users of the OSN platform, on average, react to the policy change by opening up their boundaries to be less private and/or more public in sharing, we can argue the level of reaction might be different for different groups of users. Specifically, we believe

that those who were more private in sharing before the event are to loosen their privacy boundary at a larger extent. The argument is simple: users who were more private in their disclosure behaviour before are more likely to use granular controls brought about by the privacy policy change to properly choose the audience for their posts. Hence, the policy change is expected to cause these users to use more public posts relative to private messages because audience can be customized for each public post. However, those who were not so sensitive to privacy before the event are less likely to change their disclosure patterns based on the new privacy controls. Hence, we expect that the level of reaction to the privacy policy change is stronger for users who were more private in their disclosure before the policy change. We will also test this prediction in our ongoing research. We hope to have the results from new analysis available by WEIS 2013.

Conclusions

In this study, we examine the relationship between privacy controls and content sharing activities of users in the context of an online social network. Our identification is based on the exogenous policy change in the social networking platform in December 2009 that brought about more granular privacy controls. New controls enable users to choose the intended audience for each wall post instead of using the same audience for all posts. We find that disclosure patterns of users based on content sharing activities have changed to reflect the openness in sharing as a result of improved privacy controls. Specifically, OSN users increased the use of wall posts and decreased the use of private messages. Simplifying complex privacy settings and giving users more control over their information indeed resulted in sharing content more openly. We attribute this change to the sense of control hypothesis: OSN users fear of privacy loss and the reduction in this fear results in higher levels of disclosure. Although the policy change was portrayed as a better tool for privacy protection, our results clearly show that the OSN platform achieved more open sharing of social networking activities with the help of the policy change.

We acknowledge that we use the frequency of different content sharing activities to characterize the level of openness (or closedness) in disclosure patterns of users. It is possible that OSN users generate more wall posts after the policy change and yet use the in-line controls extensively to reduce the size of the reach of their wall posts. We do not have data on how often users take advantage of new privacy controls to restrict who can see their wall posts. This is a limitation of our study. However, anecdotal evidence suggests that people seldomly use in-line controls to limit the audience for their wall posts (Kincaid 2009). In addition, as part of the policy change, the OSN

platform loosened the recommended privacy setting of users who have never edited their settings before from "only friends" to "everyone" for wall posts, leading to an increase in openness by default. Given that only 15-20% of users ever changed their privacy settings (including the setting for wall posts), the OSN platform effectively switched 80-85% of users to share their status updates, links, pictures and videos with the whole web (Kirkpatrick 2009). Furthermore, the audience selector tool remembers the audience a user shared with the last time the user posted something, and uses the same audience when the user shares again unless the user changes it. It is highly likely that users who have never tuned their privacy settings will share each and every wall post after the change with the public by default. Hence, we can argue that the decrease in privateness index we identified in this study due to the policy change indeed implies an increase in openness of content sharing on the platform.

Apart from a significant empirical finding, our study contributes to the privacy literature in new ways. First, prior privacy studies almost exclusively relied on experimental data to assess the potential impact of a change in privacy policies or privacy controls. Instead of conducting an experiment on a small group of users, our empirical results come from observational data from a large user network. Second, earlier studies often used "intention" of disclosure as a dependent variable. However, the privacy paradox proves that individuals actually disclose more than their reported intentions. In this study, we measured the "actual" and observed disclosure behaviour and showed the relationship between privacy controls and observed disclosure behaviour. Third, we introduce to the literature a simple yet a very useful metric, called privateness index, to capture the patterns of content sharing activities of OSN users in relative terms. This index serves as a proxy for openness and closedness of shared content. Fourth, our results are also consistent with recent experimental finding about control dilemma in privacy. Changes that give users more control over their personal profile information may have unintended consequence of eliciting greater disclosure of personal information (Brandimarte et al. 2010). Fifth, we study content generation and sharing patterns of OSN users in a longitudinal setup. This was not possible in prior research given the proprietary nature of the data. Instead of studying privacy in the context of (mostly) static profile elements, we gather and analyze dynamic content production and disclosure process. The last, but not the least, this paper contributes to the policy debate surrounding the effectiveness of solutions relying solely on user privacy controls for privacy protection of OSN users.

References

Acquisti, A., and Gross, R. 2006. "Imagined Communities: Awareness, Information Sharing, and Privacy on The Facebook," 6th *Workshop on Privacy Enhancing Technologies*, Cambridge, UK.

Ajzen, I. 1991. "The Theory of Planned Behavior," *Organizational Behavior and Human Decision Processes* (50), pp. 179–211.

Anderson, K. 2010. Is Facebook Eroding Privacy? Or Does Social Media Require Us to Lower Our Expectations? The Scholarly Kitchen. (May 10). Available at http://scholarlykitchen.sspnet.org/2010/05/10/is-facebook-eroding-privacy-or-does-social-media-require-us-to-lower-our-expectations/

Altman, I. 1975. *The Environment and Social Behavior: Privacy, Personal Space, Territory, and Crowding*, Monterey, CA: Brooks/Cole Pub. Co.

Aral, S., and Walker, D. 2011. Creating Social Contagion through Viral Product Design: A Randomized Trial of Peer Influence in Networks," *Management Science* (57:9), pp. 1623-1639.

Bankston, K. 2009. Facebook's New Privacy Changes: The Good, The Bad, and The Ugly. Electronic Frontier Foundation.

Brandimarte, L., Acquisti, A., and Loewenstein, G. 2010. "Privacy Concerns and Information Disclosure: An Illusion of Control Hypothesis," *Conference on Information Systems and Technology*, Austin, TX.

Caverlee, J. and Webb, S. 2008. "A Large-scale Study of Myspace: Observations and Implications for Online Social Networks," *International Conference on Weblogs and Social Media*, Seattle, WA.

Cheng, J. 2009. An Updated Guide to Facebook Privacy: December 2009 edition. arstechnica.com.

Fischhoff, B., Slovic, P., Lichtenstein, S., Read, S., and Combs, B. 1978. "How safe is safe enough? A Psychometric Study of Attitudes towards Technological Risks and Benefits," *Policy Science* (8), pp. 127-52.

Fowler, G.A. 2012. Facebook Tops Billion-User Mark. *The Wall Street Journal* (October 4).

Freeman, K. 2012. Facebook Privacy: This Service Alerts You When it Changes [INFOGRAPHIC] Mashable (May 14). Available at http://mashable.com/2012/05/14/privacywatch-infographic/

Gross, R., and Acquisti, A. 2005. "Information Revelation and Privacy in Online Social Network (The Facebook Case)," *ACM Workshop on Privacy in the Electronic Society*.

Jones, H., and Soltren, J. H. 2005. "Facebook: Threats to Privacy," December 14, Available at http://www-swiss.ai.mit.edu/6805/student-papers/fall05-papers/facebook.pdf.

Kincaid, J. 2009. The Facebook Privacy Fiasco Begins. TechCrunch (December 9). Available at http://techcrunch.com/2009/12/09/facebook-privacy/

Kirkpatrick, M. 2009. Why Facebook Changed its Privacy Strategy. ReadWrite (December 10). Available at http://readwrite.com/2009/12/10/why_facebook_changed_privacy_policies/

Lampe, C., Ellison, N. B., and Steineld, C. 2007. "A Face(book) in the Crowd: Social Searching vs. Social Browsing," 20th *Conference on Computer Supported Cooperative Work*, New York, NY

Langer, E. J. 1975. "The Illusion of Control," *Journal of Personality and Social Psychology* (32:2), pp. 311-328.

Laufen, R. S. and Wolfe, M. 1977. "Privacy as a Concept and a Social Issue - Multidimensional Developmental Theory," *Journal of Social Issues (*33), pp. 22-42.

Lewis, K., Kaufman, J., and Christakis, N. 2008. "The Taste for Privacy: An Analysis of College Student Privacy Settings in an Online Social Network," *Journal of Computer-Mediated Communication*, (14:1), pp. 79-100.

Margulis, S. T. 2003. "Privacy as a Social Issue and Behavioral Concept," *Journal of Social Issues (*59), pp. 243-261.

Naone, E. 2010. The Changing Nature of Privacy on Facebook. MIT Technology Review (May 3). Available at http://www.technologyreview.com/news/418766/the-changing-nature-of-privacy-on-facebook/

Opshal, K. 2010. Facebook's Eroding Privacy Policy: A Timeline (April 28). Available at https://www.eff.org/deeplinks/2010/04/facebook-timeline/

Peltzman, S. 1975. "The Effects of Automobile Safety Regulation," *Journal of Political Economy* (83), pp. 677–726

Petronio, S. 2002. *Boundaries of Privacy: Dialectics of Disclosure*, New York, NY: Sunny Press.

Sanghvi, R. 2009. New Tools to Control Your Experience. The Facebook Blog (December 9).

Stutzman, F. 2006. "An Evaluation of Identity-sharing Behavior in Social Network Communities," *International Digital Media and Arts Journal*, (3:1), pp.10-18.

Stutzman, F., Gross, R., and Acquisti, A. 2012. "Silent Listeners: The Evolution of Privacy and Disclosure on Facebook," *Journal of Privacy and Confidentiality* (4:2), pp. 7-41.

Thelwall, M. 2008. "Social Networks, Gender and Friending: An Analysis of Myspace Member Profiles," *Journal of the American Society for Information Science and Technology,* (59:8), pp. 1321-1330.

Thompson, S. C. 1999. "Illusions of Control: How We Overestimate Our Personal Influence," *Current Directions in Psychological Science* (8:6), pp. 187–190.

Tufekci, Z. 2008. "Can You See Me Now? Audience and Disclosure Regulation in Online Social Network Sites," *Bulletin of Science Technology and Society*, (28:1), pp. 20-36.

Williams, D. 2012. How Facebook's 'Frictionless Sharing' can Create Better Ads on Facebook. AdAge Digital. Available at http://adage.com/article/digitalnext/facebook-s-frictionless-sharing-create-ads-facebook/232419/